

A SciDAC retrospective from USQCD

B. Joó, on behalf of the USQCD Collaboration
 Thomas Jefferson National Accelerator Facility (Jefferson Lab)
 12000 Jefferson Avenue, Newport News, VA 23606

Abstract

The USQCD collaboration, has developed a coordinated software infrastructure that has allowed it to successfully exploit virtually every major computer architecture during the lifetime of the current SciDAC project. As this particular SciDAC project reaches its final year, we review some of the particular software successes of the collaboration and present highlights of scientific achievements enabled by this work. We finish with a look toward the exascale horizon.

I. INTRODUCTION

Many outstanding questions in nuclear and high energy physics require a quantitative understanding of the theory of the strong nuclear force known as quantum chromodynamics (QCD): Why is there more matter than anti-matter in the universe? What are the relevant degrees of freedom for low energy QCD, and what is the role of the gluons? How are protons and neutrons bound together into nuclei? What happens to quark matter under conditions of extreme temperature and pressure when they melt into a quark gluon plasma? These questions all form the subject of intense experimental investigation at accelerator facilities such as the Tevatron at Fermi National Accelerator Laboratory (FNAL), Jefferson Lab (JLab), and the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL). The answers are important to our fundamental understanding of the universe, for potential future energy production, and for nuclear material stewardship.

One versatile computational technique that can be used to investigate questions of QCD and other QCD-like theories is lattice gauge theory (LGT). The lattice version of QCD, known as lattice QCD (LQCD), currently provides the only known model-independent ab initio method of carrying out nonperturbative QCD calculations. By discretizing the equations of QCD onto a space-time lattice and performing a rotation to Euclidean time, the original quantum field theory is transformed to a statistical mechanical system that is then amenable to computer simulation.

The USQCD collaboration comprises nearly all lattice gauge theory practitioners in the United States. The scientific program of the collaboration ranges through all the areas of nuclear and high energy physics mentioned previously. As a SciDAC project, the USQCD collaboration has coordinated the development of software to target computer architectures ranging from leadership computing facilities and special-purpose computers for QCD to clusters with commodity processors and interconnects. USQCD also carries out algorithmic research as well as software development and optimization to extract good performance from the available hardware and to prepare for future architectures. This work is often carried out in collaboration with other SciDAC centers and institutes such as the TOPS (Towards Optimal Petascale Solvers) collaboration and the Performance Engineering Research Institute (PERI).

In the remainder of this contribution, we will take a retrospective view of some of the software and scientific achievements of the USQCD collaboration during this SciDAC project and present our outlook for the future.

The rest of this contribution is organized as follows. In Sec. II we briefly review the structure of LQCD calculations. In Sec. III we present an overview of the computer hardware that was available to the USQCD Collaboration during the lifetime of this SciDAC project and discuss the software infrastructure that has been developed to allow the successful exploitation of these architectures. In Sec. IV we present a sampling of major scientific achievements. In Sec. V we present a summary and consider future work.

II. ELEMENTS OF AN LQCD CALCULATION

Lattice QCD calculations split naturally into two phases. The first phase is called *gauge generation*, wherein ensembles of representative samples of the QCD vacuum; *gauge configurations* are generated through a Markov chain Monte Carlo method. In this step one configuration is generated from its predecessor, typically by employing a reversible, fictitious Hamiltonian molecular dynamics (MD) treating the gauge fields on the links of the lattice as canonical coordinates. The configuration at the end of the MD trajectory is subject to a Metropolis accept/reject step.

The second phase is referred to as the *analysis* phase, where observables are computed on the configurations in the previously generated ensembles, resulting in the computation of *correlation functions* that contain the physical information in the theory. In the study of hadronic matter, the key components of observables are *quark propagators*.

In both the gauge generation and analysis phases the most computationally demanding part of the calculation comes from solving the lattice Dirac equation that describes the interaction of quarks with gluons. The lattice Dirac equation is large and sparse, with a dimensionality of $O(10^8)$ in current calculations. The condition of the lattice Dirac equation is affected by the physical volume, the lattice spacing, and the quark mass. As the lattice spacing is taken toward the continuum and the quark masses approach those of the physical quarks, the condition of the lattice Dirac operator and the correlation length of the theory both diverge, leading to *critical slowing* in simulation algorithms.

III. SOFTWARE INFRASTRUCTURE

A. Hardware

As we entered the current SciDAC project, the primary facilities available for LQCD calculations in the U.S. were the QCDOC [1] supercomputer installation at BNL and cluster facilities at FNAL and JLab, which were procured and operated through the recently established national computational facility for lattice QCD. Gauge generation programs ran on the QCDOC, and analysis work was run on the cluster facilities. During the period of the SciDAC project the U.S. Department of Energy built up the leadership computing facilities at Argonne and Oak Ridge national laboratories and made available an unprecedented amount of computer time through the DOE INCITE program. USQCD competed successfully in INCITE and moved large parts of its gauge generation programmes to leadership facilities such as the Blue Gene/P system (Intrepid) at ANL and the Cray XT system (Jaguar) at ORNL. Further developments in the area of high-performance computing included the rise of multicore and heterogeneous accelerated architectures. Multisocket, multicore systems with 4–8 cores per socket have become commonplace. Heterogeneity has become mainstream through the rise of general-Purpose graphics processor (GPU) technology.

B. Layered Framework

The key to the successful exploitation of architectures ranging from clusters to Cray and Blue Gene systems and, most recently, to Nvidia GPUs has been a layered architecture. Four primary layers have been defined. The lowest layer (Level 1) provided an API for a simplified message-passing interface called QCD Message Passing (QMP) and for some basic linear algebra common in QCD (QLA).

The second level was a data-parallel layer referred to as QCD Data Parallel (QDP). This layer provided operations that would be carried out uniformly on the whole lattice, such as additions of lattice vectors, shifts, and lattice-wide reductions. This layer was intended primarily as a productivity layer.

The third layer is an optimization layer referred to as Level 3. This layer contains machine-specific optimizations that can cut through Level 2 and get close to hardware, should the data parallel layer API not provide sufficient performance. Level 3 is in some sense notional and is implemented as a variety of modules that can often be installed in isolation. Examples of Level 3 modules include the QDPQOP package, the MDWF solver for domain wall fermions [2], and most recently the QUDA library for GPUs [3], [4]

The top layer comprises the main application frameworks. Three large frameworks are in use in USQCD: The MILC code [5], the Chroma code [6], and the Columbia Physics System (CPS) [7]. Both MILC and CPS predate the layered framework and integrated the layers defined above on an as-needed basis to meet portability and performance aims. Chroma, in contrast, has been built entirely in terms of these layers, mostly in terms of the Level 2 QDP++.

C. Algorithmic Advances

Algorithmic advances in LQCD have fit into four primary areas: improved theoretical formulations of LQCD especially in the areas of Chiral Fermions [8], [9], [10]; improved linear solvers [11], [12], [13], [14], [3]; improved molecular dynamics algorithms [15], [16]; and improved analysis techniques [17], [18]. Because of lack of space we will concentrate mostly on the area of solvers in this section.

As the quark mass is decreased and the lattice spacing is made fine, the critical slowing in the solver is usually caused by the gap in the near-zero part of the eigenspectrum “closing.” In other words, the lowest eigenvalue of the linear operator creeps closer to zero, and the density of low modes increases. One way to deal with this situation, if the spectrum is not yet very dense, is to explicitly deflate the low modes as is done in the EigCG and EigBiCG algorithms [11], [12].

A second successful way, developed through a collaboration of USQCD and the TOPS, is through the use of adaptive, algebraic, multigrid techniques [13], [14]. In Figure 1 we reproduce a plot from [14] illustrating the behavior of the new multigrid approach. It can be seen that critical slowing is completely eliminated as compared with the BiCGStab solver.

Both the deflated and multigrid solvers have a start-up cost, which requires several solves to amortize. In a large-scale analysis project, where one needs to solve the system on the same configuration but for many right-hand sides, this setup cost is negligible. The use of multigrid and deflated solvers in gauge generation is an open area of research.

Recently the QUDA library of solvers has been developed for leveraging the power of CUDA-enabled GPUs [3], [4]. On GPU systems, memory bandwidth is a serious bottleneck for QCD calculations, which in single precision typically require close to 1 flop/byte. The QUDA library developed a variety of algorithmic improvements to fight the memory bandwidth restrictions including the use of multiple-precision solvers, on-the-fly (de)compression of the fields used in the calculation by appealing to the symmetries of the theory, and the merging of independent BLAS operations to aid data reuse. Many of these improvements have been back-ported for non-GPU use, for example in the Chroma code. We show in Fig. 2 the scaling of the multiprecision BiCGStab solver implementation on a variety of clusters from [19]. Currently the parallelization in QUDA has been extended to multiple dimensions; and through the use of a domain decomposed solver, the library has been successfully strongly scaled to 256 GPUs on the Edge Cluster at LLNL, details of these latest advances will be reported elsewhere [20].

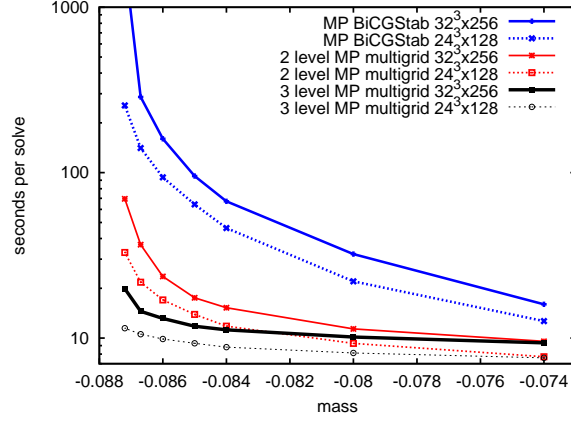


Fig. 1: Performance of multigrid solvers compared with BiCGStab (from [14]).

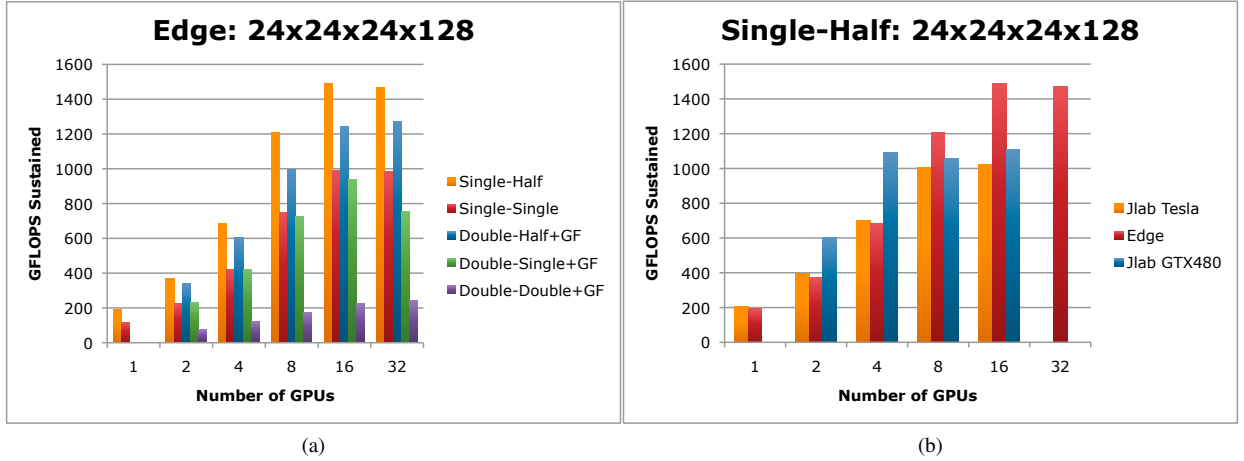


Fig. 2: Performance of BiCGStab Solver on multiple GPUs on JLab GPU clusters and the Edge Cluster. Left: varying precision on Edge. Right: Performance effects of different cluster architectures. Edge has 2 Tesla M2050 GPUs per node and full QDR Infiniband (IB), while the JLab clusters have 4 GPUs per node (Tesla or GTX480) and have QDR IB cards, sitting in half-width (x8) slots. From [19].

IV. SCIENTIFIC HIGHLIGHTS

We finish our review with two scientific highlights, one taken from high energy physics and one from nuclear physics. Further nuclear physics results will be discussed in the plenary contribution to this conference by Savage.

In Figure 3 we show the current constraints on a Unitarity Triangle (UT) from [21], which combines data from LQCD calculations of the bag parameter B_K [22], [23], [24], [25] (light purple band marked $\epsilon_K + |V_{cb}|$) and from LQCD calculations of $B - \bar{B}$ mixing [26], [27], [28] (green band, labeled $\frac{\Delta M_s}{\Delta M_d}$).

This UT is a graphical depiction of the constraints on the CP violating parameters ρ and η in the Standard Model. The area of the triangle is a measure of CP-violation, which is related to the abundance of matter over antimatter in the universe.

The constraints coming from the lattice calculations (light purple band, green band) form a broad intersection, which barely intersects the constraint from the $S\psi K$ cone (bright purple). The possibility that a discrepancy between the constraints may arise, should the precision of the lattice calculations improve, is an intriguing one, as it may point to a chink in the armour of the Standard Model, and potentially to new physics beyond. Pushing the precision and accuracy of the lattice QCD calculations to the same order as experiment to resolve this issue is a major science driver for the USQCD collaboration.

In Figure 4 we show a recent determination of the spectrum of isoscalar mesons from the Hadron Spectrum collaboration [29]. The height of the boxes on the figure is the statistical error on the determination of the energy of the states, whereas the relative area of black to green color in the boxes shows a measure of the light quark to strange quark content of the mesons. Grey boxes indicate corresponding isovector states [30], and the pink boxes indicate the masses of glue-ball states measured in a quark free Yang Mills theory [31]. We note the determination of several exotic mass states to high accuracy, and at energies which should be accessible to JLab after the 12 GeV upgrade. This calculation had to surmount three major difficulties: First,

Figure 1 is a plot showing the ratio of physical to constituent quark masses, $m_{\text{phys.}} / m_Q$, for various meson states. The y-axis ranges from 0.5 to 2.5 GeV. The x-axis is divided into three regions: negative parity, positive parity, and exotics. The legend indicates isoscalar (black/green), isovector (gray), and YM glueball (pink) states. The plot shows that for many states, the physical mass is significantly higher than the constituent quark mass, especially for exotic states and higher angular momentum states.

excited states are difficult to extract, because the signal dies away very rapidly. Second, spin classification is difficult since the hypercubic lattice on which the theory is formulated breaks rotational symmetries. Third, isoscalar states receive contributions from disconnected diagrams, which are statistically very noisy. In order to surmount these challenges, anisotropic lattices were used whose fine temporal extent aids in resolving the rapidly falling signals of the excited states. The spin identification required a large basis of operators and a variational calculation, while dealing with the noise from disconnected diagrams required a whole new algorithmic approach to the analysis through the use of a technique called distillation [17]. Hence this result is a culmination of algorithmic and code developments that were all developed under SciDAC and a gauge generation project that has been enabled by the availability of time through the INCITE program.

The past five years have been witness to many advancements in the area of high-performance computing and have resulted in impressive scientific calculations through the use of lattice gauge theory. We have mentioned but a few of the computational and scientific achievements of the USQCD collaboration in this contribution. Going forward, it is clear that substantial algorithmic and software investments will be required in order to adapt LQCD software to the characteristic features of future architectures on the road to the exascale.

The distortion of the balance between floating point capability, and the latency and bandwidth of local and remote memory accesses will become the key factor which determines the suitability of particular algorithms for a given architecture. As this imbalance becomes more and more severe, substantial research into new solvers and other algorithms will be required.

High-level, domain-specific frameworks such as Chroma can be a boon in this situation, since the solvers they wrap can be swapped in and out as needed depending on the architectures. At the same time, lower levels of the software layers on which the framework depends (QDP++ in the case of Chroma) will also need to be ported and optimized on the new architectures to avoid Amdahl's law bottlenecks. Such porting and optimization may require substantial re-engineering to cope with new programming models; for example, to deal with heterogeneity, or algorithmic requirements that may not have been present in the design of the original software.

Surmounting this challenge will require collaboration with applied mathematicians, computer scientists, LQCD domain specialists and hardware vendors. We look forward to an exciting and productive follow-on SciDAC project.

ACKNOWLEDGMENTS

B. Joó gratefully acknowledges funding through U.S. DOE project grants DE-FC02-06ER41440 and DE-FC02-06ER41449 (USQCD SciDAC project) and DE-AC05-06OR23177 under which Jefferson Science Associates LLC manages and operates Jefferson Lab. B. Joó would like to thank Ruth van de Water for assistance on the discussion regarding the Unitarity Triangle and for contributing the graphic from [21], and Robert Sugar and Ruth van de Water for careful reading of, and helpful comments this manuscript. The U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce this manuscript for U.S. Government purposes.

REFERENCES

- [1] P. Boyle, D. Chen, N. Christ, M. Clark, S. Cohen, Z. Dong, A. Gara, B. Joo, C. Jung, L. Levkova, X. Liao, G. Liu, R. Mawhinney, S. Ohta, K. Petrov, T. Wettig, A. Yamaguchi, and C. Cristian, "QCDOC: A 10 Teraflops Computer for Tightly-Coupled Calculations," in *Proceedings of the ACM/IEEE SC2004 Conference*, ser. SC '04, 2004, p. 40.
- [2] A. Pochinsky, "Mdwf moebius domain wall solver implementation," <http://www.mit.edu/~avp/mdwfl/>.
- [3] M. A. Clark, R. Babich, K. Barros, R. C. Brower, and C. Rebbi, "Solving Lattice QCD systems of equations using mixed precision solvers on GPUs," *Comput. Phys. Commun.*, vol. 181, pp. 1517–1528, 2010.
- [4] R. Babich, M. A. Clark, and B. Joó, "Parallelizing the QUDA Library for Multi-GPU Calculations in Lattice Quantum Chromodynamics," in *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 1–11. [Online]. Available: <http://dx.doi.org/10.1109/SC.2010.40>
- [5] M. L. Collaboration, C. Bernard *et al.*, "The milc code," <http://www.physics.utah.edu/~detar/milc/milcv7.pdf>, 2010.
- [6] R. G. Edwards and B. Joo, "The Chroma software system for lattice QCD," *Nucl. Phys. Proc. Suppl.*, vol. 140, p. 832, 2005.
- [7] Riken-Brookhaven-Columbia (RBC) Collaboration, "The Columbia Physics System (CPS)," http://phys.columbia.edu/~cqft/physics_sfw/physics_sfw.htm, 2010.
- [8] R. C. Brower, H. Neff, and K. Orginos, "Moebius fermions," *Nucl. Phys. Proc. Suppl.*, vol. 153, pp. 191–198, 2006.
- [9] R. G. Edwards, B. Joo, A. D. Kennedy, K. Orginos, and U. Wenger, "Comparison of chiral fermion methods," *PoS*, vol. LAT2005, p. 146, 2006.
- [10] P. M. Vranas, "Gap domain wall fermions," *Phys. Rev.*, vol. D74, p. 034512, 2006.
- [11] A. Stathopoulos and K. Orginos, "Computing and deflating eigenvalues while solving multiple right hand side linear systems in Quantum Chromodynamics," *SIAM J. Sci. Comput.*, vol. 32, no. 1, pp. 439–462, 2010.
- [12] A. Abdel-Rehim, K. Orginos, and A. Stathopoulos, "Extending the eigCG algorithm to non-symmetric linear systems with multiple right-hand sides," *PoS*, vol. LAT2009, p. 036, 2009.
- [13] R. Babich *et al.*, "Adaptive multigrid algorithm for the lattice Wilson-Dirac operator," *Phys. Rev. Lett.*, vol. 105, p. 201602, 2010.
- [14] J. C. Osborn *et al.*, "Multigrid solver for clover fermions," *PoS*, vol. LATTICE2010, p. 037, 2010.
- [15] M. A. Clark, A. D. Kennedy, and P. J. Silva, "Tuning HMC using Poisson brackets," *PoS*, vol. LATTICE2008, p. 041, 2008.
- [16] A. Kennedy, M. Clark, and P. Silva, "Force Gradient Integrators," *PoS*, vol. LAT2009, p. 021, 2009.
- [17] M. Peardon *et al.*, "A Novel quark-field creation operator construction for hadronic physics in lattice QCD," *Phys.Rev.*, vol. D80, p. 054506, 2009.
- [18] C. Morningstar *et al.*, "Excited-State Hadrons using the Stochastic LapH Method," 2011.
- [19] Joo, B. and Clark, M. A., "Lattice QCD on GPU Clusters, using the QUDA Library and the Chroma Software System," talk given at ICCS Workshop, Berkeley CA, Jan 2011, submitted for publication in International Journal Of High Performance Computing Applications, 2011.
- [20] Babich, R. and Clark, M. A. and Shi, G. and Joo, B. and Brower R. and Gottlieb S., "Scaling Lattice QCD beyond 100 GPUs," Submitted for publication to the International Conference for High Performance Computing, Networking, Storage and Analysis, SC'11 to be held in Seattle, Nov 2011.
- [21] J. Laiho, E. Lunghi, and R. S. Van de Water, "Lattice QCD inputs to the CKM unitarity triangle analysis," *Phys.Rev.*, vol. D81, p. 034503, 2010.
- [22] Y. Aoki, R. Arthur, T. Blum, P. Boyle, D. Brommel *et al.*, "Continuum Limit of B_K from 2+1 Flavor Domain Wall QCD," 2010.
- [23] C. Aubin, J. Laiho, and R. S. Van de Water, "The Neutral kaon mixing parameter $B(K)$ from unquenched mixed-action lattice QCD," *Phys.Rev.*, vol. D81, p. 014507, 2010.
- [24] T. Bae, Y.-C. Jang, C. Jung, H.-J. Kim, J. Kim *et al.*, " B_K using HYP-smear staggered fermions in $N_f = 2 + 1$ unquenched QCD," *Phys.Rev.*, vol. D82, p. 114509, 2010.
- [25] E. Gamiz *et al.*, "Unquenched determination of the kaon parameter $B(K)$ from improved staggered fermions," *Phys.Rev.*, vol. D73, p. 114502, 2006.
- [26] E. Gamiz, C. T. Davies, G. Lepage, J. Shigemitsu, and M. Wingate, "Neutral B Meson Mixing in Unquenched Lattice QCD," *Phys.Rev.*, vol. D80, p. 014503, 2009.
- [27] R. Evans, A. El-Khadra, and E. Gamiz, "A determination of the $B_0(s)$ and $B_0(d)$ mixing matrix elements in 2+1 lattice QCD," *PoS*, vol. LATTICE2008, p. 052, 2008.
- [28] C. Albertus, Y. Aoki, P. Boyle, N. Christ, T. Dumitrescu *et al.*, "Neutral B-meson mixing from unquenched lattice QCD with domain-wall light quarks and static b-quarks," *Phys.Rev.*, vol. D82, p. 014505, 2010.
- [29] J. J. Dudek *et al.*, "Isoscalar meson spectroscopy from lattice QCD," *accepted for publication as a Rapid Communication in Phys. Rev. D. arXiv:1102.4299 [hep-lat]*, 2011.
- [30] —, "Toward the excited meson spectrum of dynamical QCD," *Phys. Rev.*, vol. D82, p. 034508, 2010.
- [31] C. J. Morningstar and M. J. Peardon, "The glueball spectrum from an anisotropic lattice study," *Phys. Rev.*, vol. D60, p. 034509, 1999.